PERGAMON

Telecommunications Policy 24 (2000) 113–134

TELECOMMUNICATIONS
POLICY

www.elsevier.com/locate/telpol

The networks of the Internet: an analysis of provider networks in the USA

Sean P. Gorman, Edward J. Malecki*

Department of Geography, University of Florida, PO Box 117315, Gainesville, FL 32611-7315, USA

Abstract

The Internet is comprised of a large number of private and public networks, which function as autonomous systems within the overall network. The network structure of an individual network, which may have grown by acquisition, has an impact on its efficiency. Using graph theory, this paper examines the network structure of ten "backbone" provider networks in the USA. The networks exhibit very different structures, which affect both the competitive positions of the firms as providers of Internet service to their customers and the nature of integration of other networks.

Keywords: Internet; Network; Network analysis; Graph theory

1. Introduction

The Internet is a remarkable convergence or fusion of information types, information media, and information operators (Kellerman, 1997). The convergence of industries has changed the products, functions, technologies, and services provided by a soft-network of applications that remain dependent on access to a hard-network infrastructure (Antonelli, 1997). The Internet already has been identified as a *general-purpose technology*, joining writing, printing, electricity, and a handful of more recent technologies, including lasers, the factory system, mass production, and flexible manufacturing (Lipsey, Bekar & Carlaw, 1998).

Although the Internet is a network of networks, it has been the subject of surprisingly little network analysis. Four examples of network analysis illustrate the versatility of the method, but also the small amount of research to date to probe the networks that comprise the "cloud" with

* Corresponding author. Tel.: + 1-352-392-0496; fax: + 1-352-392-8855.
E-mail address: malecki@geog.ufl.edu (E. J. Malecki)

PII: S0308-5961(00)00005-7

which the Internet is commonly portrayed.[1] First, Wan and Chung (1998) have utilized network analysis to capture the structure of links among Web pages within a Web site. They treat the time required to download or transmit a page as the cost of a link. In a second example, Burch and Cheswick (1999) have attempted to map the entire Internet. Tracing routes to IP addresses to routers allowed them to create a "map" of the Internet, but they were unable to display the 88,000 unique IP addresses, or nodes, in their data base (of approximately 94,000 networks connected to the Internet). They make in passing a key observation for a network analysis: "as far as the Internet is concerned, the fact that Carnegie Mellon University and Lycos are 14 hops away is more important than that they are four miles apart" (Burch & Cheswick, 1999, p. 98). The logical layout they describe in terms of *hops* — based on connectivity rather than location — is what characterizes network analysis.[2] In a third example, "Internet tomography" — analogous to a medical CAT scan — depicts the topology of the Internet based on infrastructure-wide (global) connectivity and round-trip time for packets sent along a path from A to B (Claffy, Monk & McRobb, 1999). Fourth, Wheeler and O'Kelly (1999) have used network analysis to identify the network topology of 31 US Internet backbones and the Internet accessibility of 141 cities. This paper offers an improvement over the methodology used by Wheeler and O'Kelly.

Before moving to an analysis of the networks of the major Internet backbone providers in the USA, we provide a brief history of the Internet and its backbone networks, and the evolution of the corporate structure of the Internet. The paper then introduces network analysis and its application to ten Internet backbones in the USA. The paper concludes with some implications of this analysis for understanding the Internet.

2. A brief history of the Internet

The Internet is a worldwide network of computer networks that use a common communications protocol, Transmission Control Protocol/Internet Protocol (TCP/IP). It began in the late 1960s as ARPANET, a project of the Advanced Research Projects Administration of the US Defense Department, designed to link together universities and defense contractors using TCP/IP packet-switching technology. In the mid-1980s the National Science Foundation (NSF) extended the technology and, through NSFNET, took the "leadership in networking" (Rogers, 1998). NSFNET provided connectivity to NSF's supercomputer centers and provided a high-speed backbone for developing the Internet (Mackie-Mason & Varian, 1997).

NSFNET's "acceptable use" policies allowed only academic and research locations to be linked and send information on NSFNET backbone. Corporations wanting to interconnect their corporate systems with TCP/IP contracted with private telecommunications providers (Schiller, 1999). To meet the growing market demand for similar communications systems in the private market, three firms (Alternet, PSInet and SprintLink) formed backbones of their own as private, for-profit

[1] The "cloud" may be the entire Internet, an analogy for its lack of fixed form (Bailey, 1997), or it may symbolize the thousands of autonomous systems (ASs) which comprise the Internet (Srinagesh, 1997).

[2] Other efforts to map cyberspace in its many dimensions, rather than only as a network, are catalogued in the Atlas of Cyberspace < http://www.cybergeography.org/atlas/atlas.html >

enterprises (Mackie-Mason & Varian, 1997). MCI, which was contracted to run NSFNET, had its early lead solidified into a leading market share today. In April 1995, NSFNET was shut down, which left for-profit enterprises running the new backbone infrastructure at the heart of the Internet. Those firms today largely determine Internet policy, interconnectivity and usage. Most technological decisions are made by small committees of volunteers, such as the Internet Engineering Task Force (IETF) (Gillett & Kapor, 1997). The Internet Corporation for Assigned Names and Numbers (ICANN) is the new non-profit corporation that was formed to take over responsibility for the IP address space allocation, protocol parameter assignment, domain name system management, and root server system management functions now performed under US Government contract by the Internet Assigned Numbers Authority (IANA) and other entities (ICANN, 1999).

With commercialization came the global spread of the Internet. Internet globalization, however, has a peculiar structure, because the Internet still revolves — physically, technologically and economically — around its origin, the USA. It is not a uniform global network, but a centralized network radiating from the United States. Greater bandwidth (high-speeds links) to the US than between locations within countries results in an Ameri-centric network (Cukier, 1999; Evagora, 1997; Paltridge, 1998a).

Deregulation has moved the vast majority of the world's telecommunication networks into private control. Even before deregulation, private firms had supplied advanced telecommunication services (Gabel, 1996). Privately operated networks in a competitive environment respond mainly to market pressures of supply and demand (Graham & Marvin, 1996; Salomon, 1996). Where demand is greatest, telecommunications will be supplied, resulting in geographic biases. The principal bias is that agglomerations of population and economic activity are supplied telecommunications infrastructure and services disproportionately. The agglomeration of demand in urban areas makes it most unlikely that the Internet and related technologies will result in the "end of geography," the "death of distance" and the decentralization of economic activity (Gillespie & Robins, 1989; Moss, 1998).

The rest of this paper attempts to sketch the network structure of the new "electronic geography." The next section presents the corporate structure of the Internet and its interconnections to provide a context for the geographic implications.

3. Corporate structure of the Internet

Although the Internet often is depicted as an amorphous cloud, it has a physical structure and hierarchy. As a network of networks, the Internet is composed of a wide variety of small, medium and large networks that interconnect to give the impression to the user of one seamless machine. The networks that comprise the Internet are commonly called Internet service providers (ISPs), which vary in size as well as function.

Cukier (1998) proposes a functional classification of ISPs based on four classes: (1) transit backbone ISPs; (2) downstream ISPs; (3) online service providers such as America Online or CompuServe; and (4) firms that specialize in Web site hosting, such as Exodus. All networks below the *transit backbones* depend on the backbone ISPs either to furnish Internet connectivity or to manage the network infrastructure (Cukier, 1998). This places transit backbones at the top of the Internet "food chain" as the networks that provide the transit services that make the Internet

global. The emergence of backbone predominance is a recent phenomenon; in the openly interconnected Internet, all networks were considered more or less equals (Rogers, 1998).

Downstream ISPs primarily include hundreds of local and regional ISPs that serve individuals and small and medium-sized businesses. Because of their dependence on the backbone ISPs, downstream ISPs have been shifted by the backbone providers from network equals to paying customers (Cukier, 1998, p. 125). This cost-prohibitive dependence on backbone ISPs has forced many downstream ISPs out of business or to be acquired as backbone ISPs add to their customer base.

Web hosting companies, by contrast, have been growing, based on the value-added services they provide for customers' Web sites, but they have not escaped dependence on the backbone ISPs. Indeed, several backbone firms, such as Frontier GlobalCenter, GTE Internetworking, and Verio, also provide Web hosting, often following acquisitions of other firms (*Boardwatch*, 1999). Web hosting companies do exactly what their name implies: they host or house Web sites so they can be easily accessed by the Internet public. The problem is that the stable of Web sites connected by a Web hosting ISP create almost entirely unidirectional traffic: "The few bits of data that trickle in when a user requests a Web page are overwhelmed by the flood of outgoing audio, video, images and text" (Cukier, 1998, p. 123). This becomes problematic when it comes to deciding who is going to pay to transit this high-capacity traffic.

Online service providers, another growing segment of the market, create revenue by packaging the Internet into a user-friendly format that focuses on content and ease of use. America Online has used this format to become one of the most successful Internet firms in the market. The fact that AOL uses a proprietary protocol (i.e. other than TCP/IP) for dial-up access distinguishes it from other ISPs. Online service providers are less concerned with transmission capacity and connectivity than with gaining individual customers, predominantly residential users who connect through dial-up analog modems. As a result, online service providers lease connectivity through backbone ISPs and only manage points of presence (POPs) that connect dial-up customers to the backbone.

The national backbone and transit level network connects city nodes and provides transport for data across long-haul geographic distance. Forty private providers currently operate the national backbone transit layer. Each backbone provider is called, in network terms, an *autonomous system*, or AS. Each AS operates independently of the others, setting its own policies and determining its network structure. Only through interconnection or peering do these networks become interconnected to form the Internet.

At the technological core of the Internet are the routers, the nodes that direct traffic on the Internet, receive packets of data sent over the network and decide where the next hop will take them on the way to their destination. To know where to send a packet to get it closer to its destination, the router must know where other routers on the network are located. Even in a small network the routing tables are fairly long and complex. Ideally, a router would maintain an entry for every router of the Internet in its routing tables. As the Internet grew, routers simply could not store all the data needed to route efficiently, maintaining instead only a limited subset of the Internet with precise routing information. "This subset is not based on proximity, but economics. Business partners will have routing entries while non-economic interests are given default routes to reach other networks" (Huitema, 1997, p. 65). Routing tables represent more than a technical issue; they also embody business relationships between backbone providers.

IPv6, the new version of Internet protocol, was designed to incorporate a new hierarchy based on provider networks. Instead of listing every router in the network, a routing table would need only one entry per network (or AS), and packets would then diffuse through that provider network to their destination. Significantly, this puts network providers in control of how and to whom traffic is routed (Huitema, 1997).

Networks do not always follow political or geographical boundaries, particularly for international traffic. Internal European connections are typically routed through Washington, DC or New York. The Internet remains efficient while intra-European traffic is routed through the United States because of the relationship between *bandwidth, distance* and *delay* on IP networks. For Internet performance and efficiency, the amount of bandwidth between two places is far more important than the distance between them. In Asia, for instance, it is not uncommon for internal Internet traffic to be quicker if it is routed via California rather than over a direct link (Table 1).

Bandwidth tells only part of the story as to why the Internet's routing seems so circuitous and prefers routing through the United States over direct routes within Europe and Asia. Regulation and the lack of telecommunication competition make it more expensive to operate through Asian and European providers (Bond, 1997, p. 4; Cukier, 1998, 1999). The high cost of infrastructure and connections in Europe makes a circuit from Washington, DC to Paris, London, or Stockholm cost less than a direct line from one of these European cities to another (Paltridge, 1998a). Although prices are dropping as competition increases, leasing capacity on many intra-European leased lines remains more expensive than trans-Atlantic routes (Paltridge, 1999). These economic and technological peculiarities of the Internet have a profound impact on its network geography. More than half of intra-European and intra-Asia traffic is routed through the United States (Cukier, 1999; Evagora, 1997). The result has been an Ameri-centric network, a hub-and-spoke system, with the United States serving as the Internet hub for the world.

3.1. Peering

Backbones providers play a dominant role in the Internet, and how the backbones *peer*, or transfer data to one another is critical to its operation. Peering is more than simply transferring data; it also is allowing a peer access to a network's topology and routing tables, integrating the two networks so that data can be transited back and forth seamlessly. Originally, peering was to be done mainly at network access points (NAPs). But the growth of the Internet caused the NAPs to

Table 1
Packet round trips from Singapore Telecom Internet Exchange (STIX)[a]

Location	Distance (km)	Delay (milliseconds)	Bandwidth (Kbps)
Phnom Penh	1,100	1,100	64
Jahor Bahru	30	775	128
San Francisco	13,750	400	4,096
Vancouver	13,000	350	15,360

[a]Source: Based on Paltridge (1998a, p. 43) and Cukier (1998, p. 126).

be overburdened, resulting in a 20–30% packet loss rate (Cukier, 1998, p. 6). Those packets must then be sent again, causing further congestion of both the network and the NAPs and delay for the users (Paxson, 1997). Although the NAPs theoretically provide complete interconnectivity of the Internet, they are both public and congested.

To provide the quality of service demanded by their business customers, large backbone providers started peering privately with each other. In private peering, two or more backbone providers create a peering point to share data and routing tables only between their networks. *Peer-to-peer bilateral* interconnections are private peering points established between large firms that see themselves as equals (thus the term peers) (Bailey, 1997). Private peering has become so common that many backbone providers are leaving the NAPs entirely and are refusing to peer with smaller network providers. In order for these small companies to get their data to a non-peering provider, they must pay transit fees to stay connected. The two-party contracts define a *hierarchical bilateral* interconnection, currently the most pervasive interconnection model in today's Internet. In general, however, the large networks do not make public their peering criteria under non-disclosure agreements — nor are they required to — keeping smaller ISP's at a disadvantage (Bailey, 1997). The issue of payment, or *settlement*, for the use of transit infrastructure remains unresolved (Huston, 1999).

Hierarchical peering acknowledges the power wielded by the backbone providers. In early 1997, UUNET informed 14 ISPs that their peering agreements would be terminated, and that new bilateral transit agreements must be struck, or they would be disconnected from the Internet (Cukier, 1998). The four largest networks (UUNET, MCI, BBN [recently bought by GTE] and Sprint) control 85–95% of traffic on the Internet. Therefore, when UUNET threatened disconnection from the Internet, it was considered a serious threat. As a result of such actions, there has been discussion of possible federal regulation of the peering process to avoid the monopoly power that backbone providers are able to wield (Farrell & Katz, 1998; Rickard, 1998). At present, however, the Internet remains largely unregulated (Leo & Huber, 1997).

The Internet's structure depends critically on peering, routing tables, and provider backbones. The following section focuses on the backbones, the provider networks that comprise the Internet, and introduces network analysis of these networks.

4. Network analysis

The connection of geographically separated locations or nodes (that posses the appropriate technology levels) via discrete links is best analyzed by means of *network analysis*. Network analysis is a well-established branch of mathematics (Harary, Norman & Cartwright, 1965). Analyses of transportation and telephone networks were common during the 1960s, and remain a standard part of transportation geography (Haggett & Chorley, 1969; Taaffe, Gauthier & O'Kelly, 1996). Network analysis is well suited to the study of flows in communication networks (Nyusten & Dacey, 1961). Although analysts of telecommunications utilize the terminology of networks, sometimes little more is presented than the topology of the nodes and links within a corporate network (Hepworth, 1990; Kellerman, 1993). Thus, while the concept of a network remains central to the study of the geography of telecommunications, network analysis itself has been largely absent from recent accounts.

Transit backbones are the best choice for analysis of the structure of the Internet. Huitema (1995) suggests that the national transit backbone providers and their city nodes are the best indicators of the geography of the Internet. The backbone provider maps found in *Boardwatch's* Winter 1998 *Directory of Internet Service Providers* are used for the provider network comparisons in this research (Boardwatch, 1998). Using the *Boardwatch* network maps, cities were set as nodes and fiber optic cables of each backbone as links. Nodes within the same metropolitan statistical area (MSA) or consolidated metropolitan statistical area (CMSA) were aggregated. This is justified by the prevalence of metropolitan fiber rings and interconnected metropolitan networks functioning in the majority of large metropolitan areas.[3] This aggregation resulted in 60 nodes. A recent analysis by Wheeler and O'Kelly (1999) uses cities as nodes without aggregation, treating Palo Alto, Santa Clara and San Jose, for example, as separate nodes. Our approach aggregates these into the San Francisco–Oakland–San Jose CMSA.

The most elementary of network measures are used to compare provider networks within the Internet. The basic graph-theoretic measures of gross characteristics in Table 2A are based only on the number of nodes and links. For a more detailed look into the nature and geographic characteristics of the Internet as a network, a *connectivity* matrix must be constructed and shortest path characteristics analyzed (Table 2B). Putting the Internet into a matrix reduces it to a topological graph (Harary, Norman & Cartwright, 1965), which is especially appropriate, since distance is essentially irrelevant on the Internet. Utilizing the city/fiber optic link construct outlined above, when there is a link between two cities, a 1 is placed in the matrix; if there is no connection, a 0 is placed in the matrix. This forms a *binary connectivity matrix* of the Internet, which can be "modeled as a general graph whose nodes are AS's and whose edges are connections between pairs of AS's" (Baccala, 1998).

4.1. The network structure of the US Internet

The overall Internet infrastructure within the United States illustrates the calibration of basic graph-theoretic measures (Table 3). For the US, the number of edges (links) connecting 60 nodes was 970, just over half of the 1770 possible links. The *cyclomatic number* gives a basic indication of the size of a network. Because of the hub-and-spoke configuration of the Internet, the US *cyclomatic number* of 910 is very large.

The *beta* index indicates the complexity of a network. The USA *beta* value is extremely high for a network in general. A rail or highway network will never attain a *beta* index of greater than 3 because of its planar nature, where every intersection results in a node (vertex). A telecommunications network is non-planar, similar to an airline network: two edges can cross without the formation of a node (vertex). Theoretically, the *beta* index of a non-planar graph approaches infinity. High correlations between the *beta* indexes of national railroad networks and national development levels are common (Kansky, 1963).

A more useful index of connectivity is the *alpha* index or "redundancy index". This is the ratio between the observed number of circuits (loops) to the maximum number of circuits that could exist in a network (Haggett & Chorley, 1969). An *alpha* value of 0 would indicate a branching

[3] For definitions of MSAs and CMSAs, see http://www.census.gov/population/www/estimates/metrodef.html

Table 2
Basic graph-theoretic measures for the Internet[a]

Measure	Definition
A. Measures based on gross characteristics	
Cyclomatic number $= E - V + G$	E = number of links (edges) in the network V = number of nodes (vertices) in the network G = number of sub-graphs
Beta index $= E/V$	
Alpha index $= \dfrac{E - V + G}{\dfrac{V(V-1)}{2} - (V-1)} \times 100$	
Gamma index $= \dfrac{2E}{V(V-1)} \times 100$	
B. Measures based on shortest-path characteristics	
Diameter = maximum D_{ij}	D_{ij} = shortest path (in links) between the ith and jth node
Accessibility index $= \sum_{i=1}^{v} D_{ij}$	
Dispersion index $= \sum_{i=1}^{v} \sum_{j=1}^{v} D_{ij}$	

[a]Source: Based on Haggett and Chorley (1969, p. 32).

Table 3
Graph theoretic measures for the US portion of the Internet

Measure	USA Internet
Cyclomatic number	910
Beta index	16.167
Alpha index	51%
Gamma index	55%

network, where the removal of any one link would break the network into two sub-graphs. A value of 1 or 100% indicates a fully connected network. Values between 0% and 100% indicate the amount of redundancy in a network. For the US Internet, the *alpha* index indicates 51% redundancy, which suggests that the US Internet still has room for growth in the number of circuits connecting cities before saturation is met, although it would not be expected that every city be connected to every other city. Redundancy plays a critical role in the Internet in traffic characteristics and problems of congestion. Packets tend to use the same routing paths repeatedly over time to reach a given destination (Paxson, 1997). Therefore, building redundancy into heavily traversed routes is a key to gaining network efficiency on the Internet.

Finally, the *gamma* index is the ratio between the actual and the maximum number of edges (links) in a network. This gives an indication of the level of interconnection within the network, or the proliferation of alternate routes available to transit data from one node to another. The *gamma* index, like the *beta* index, shows a high correlation with economic development: the more

economically developed a country, the greater the number of different routes by which goods or data can be transported, increasing efficiency and decreasing congestion. The US Internet's *gamma* value of 55% suggests a high level of technological development.

5. Network analysis of backbone providers

Network analysis not only reveals the structure of the overall network of a region or country, but also can be used to examine individual networks. The US Internet "cloud" can be disaggregated into its individual proprietary networks. Each individual network depends on interconnections to others in order for the Internet to work, but end-user customers (i.e. individual homes and small businesses) generally purchase their connection through only one individual network. Graph-theoretic network analysis provides a tool for structural comparison of the individual networks that comprise the Internet network.

In a market permeated with mergers and acquisitions between competing networks, understanding the network structure of competing firms is vital and often overlooked. The recent MCI–WorldCom merger has had large impacts on both the Internet and the overall telecommunications network. The combination of MCI and WorldCom (before regulatory intervention) would have controlled over 50% of Internet backbone traffic because of structural integration of the networks of UUNET, ANS, CompuServe, and MCI (Table 4). Market dominance, along with other concerns, prompted an in-depth investigation by the US Department of Justice and the Federal Communications Commission (FCC), which ultimately approved the MCI–WorldCom merger. The primary finding of the regulators was that

> market share isn't really about backbone traffic — that's just a proxy for the real concern — the number of routes that an ISP (backbone provider) services. These downstream routes represent a network's customers where traffic either originates or terminates (Cukier, 1998, p. 130).

With this realization the US and European Community regulators added a proviso to the merger approval that MCI would have to sell its Internet backbone network before WorldCom and MCI could merge. By selling the MCI Internet backbone to Britain's Cable & Wireless (C&W), the merger was approved. (MCI WorldCom only sold the routers, servers, edge switches, multiplexers and fiber connectors to C&W, but leased the actual long-haul fiber routes to C&W, technically leaving MCI WorldCom with adjunct control of the network's fiber optic links.) Unfortunately, MCI has refused to disclose any data on the structure of its network, so a network analysis of MCI WorldCom's network with and without the MCI network is not possible.

Data on ten individual networks from the 1998 *Boardwatch Magazine's Directory of Internet Service Providers* were used to calculate graph-theoretical measures for network analysis. The *Boardwatch Directory* also included Keystone Systems' average and median download times for 50 kb test pages downloaded through connections on 15 individual networks in 26 different cities. Measurements were taken every 15 min 24 h a day for a month. The resulting download times give a measure of the efficiency of each individual network relative to each other (*Boardwatch*, 1998). Median, rather than average, download time was chosen to eliminate the effects of extreme times (in excess of 50 seconds) that skewed average download time results. The median download time gives a better view of the time that an end-user would experience the vast majority of the time, and

Table 4
IP backbone market share[a]

Network	Number of backbone connections
MCI	1888
UUNET/ANS/CompuServe	1495
Sprint	1407
GTE	354
AGIS	237
DIGEX	183
PSINET	144
CRL Comm	122
Winstar Goodnet	114
Savvis	102
Verio	93
AT&T	79
Cable & Wireless	56
DataXchange	56

[a]Source: *Boardwatch* (1998, p. 8).

is the measurement more relevant to the scope of this research. Critics of Keynote's methodology claim that the results tell more about the performance of the backbone providers' Web sites than about their actual networks (Paltridge, 1998a, p. 10). However, median download times provide an indication of typical user experience on various individual networks. The other graph-theoretic measures studied here focus on the number of routes available to an ISP, the market position indicator suggested above, rather than traffic or download speed. To look at routes, graph-theoretic network analysis must be used.

Ten individual Internet backbones were selected for graph-theoretic network analysis and comparison: UUNET, AT&T, Cable & Wireless, CompuServe, IBM, Verio, Sprint, Savvis, IDT, and Frontier. First, the traditional graph-theoretic measures of cyclomatic number, *beta* index, *alpha* index, *gamma* index, and diameter were calculated. Second, the total bandwidth available to each network (including both lit fiber and fiber devoted to IP), as it connects the 60 urban areas (MSAs or CMSAs) studied, is utilized to give a measure of the technology and gross capacity of each network. Finally, the *Boardwatch* median download time is included for a comparison of network structure to end-user experience. Table 5 gives the computations for each network in each category.

The numbers in Table 5 show a wide variation in all categories among the ten networks. The variation can be clarified when the ten networks are ranked from best to worst in each category (Table 6). The rankings in Table 6 reveal that no individual network within the Internet dominates on all measures. This is reassuring considering that the Internet depends on cooperation of many individual networks sharing resources and facilities to efficiently conduit traffic. Although there is no single dominant network, large and small networks are not equal.

Table 5
Graph theoretic measures for ten US backbone networks[a]

	UUNET	AT&T	C&W	Compu-Serve	IBM	Verio	Sprint	Savvis	IDT	Frontier
Number of edges	68	24	38	15	22	53	40	21	9	45
Number of vertices	33	16	36	11	15	26	15	19	8	10
Cyclomatic number	36	9	3	5	8	28	26	3	2	35
Beta	2.06	1.5	1.056	1.364	1.467	2.038	2.667	1.105	1.25	4.5
Alpha (%)	7.26	8.57	0.50	11.11	8.79	9.33	28.57	1.96	9.52	100.0
Gamma (%)	12.87	20.00	6.03	27.27	20.95	16.31	38.10	12.28	32.14	1.0
Diameter	5	9	9	5	5	5	4	9	4	1
Total bandwidth (Mbps)	12,558	1,080	2,675	675	990	5,372	6,806	945	405	2,025
Median download time (seconds)	2.180	2.385	3.850	2.175	3.290	2.965	2.675	2.595	1.840	2.845

[a]Source: Calculated from data in *Boardwatch* (1998).

Table 6
Backbone networks ranked on each graph theoretic measure

Edges	Vertices	Cyclomatic number	*Beta*	*Alpha*	*Gamma*	Diameter	Total bandwidth	Median download time
UUNET	C&W	Frontier (1)	Frontier	Frontier	Frontier	Frontier	UUNET	IDT
Verio	UUNET	UUNET (1)	Sprint	Sprint	Sprint	IDT (2)	Sprint	Compu
Frontier	Verio	Verio	UUNET	Verio	IDT	Sprint (2)	Verio	UUNET
Sprint	Savvis	Sprint	Verio	Compu	Compu	Compu (4)	C&W	AT&T
C&W	AT&T	AT&T	AT&T	IBM	IBM	IBM (4)	Frontier	Savvis
AT&T	IBM (6)	IBM	IBM	AT&T	AT&T	UUNET (4)	AT&T	Sprint
IBM	Sprint (6)	Compu	Compu	UUNET	Verio	Verio (4)	IBM	Frontier
Savvis	Compu	C&W (8)	IDT	IDT	UUNET	AT&T (8)	Savvis	Verio
Compu[a]	Frontier	Savvis (8)	Savvis	Savvis	Savvis	C&W (8)	Compu	IBM
IDT	IDT	IDT	C&W	C&W	C&W	Savvis (8)	IDT	C&W

[a]Compu = CompuServe. Numbers in brackets indicate ties for that rank.

In the following, we consider the ten networks from smallest to largest, generally following network size by market share as indicated in Table 4, with adaptations in this order made to take recent network acquisitions and mergers into account. For example, AT&T and IBM are discussed together because of AT&T's recent acquisition of IBM's networks, and UUNET and Sprint are discussed together because of their common links on the issue of network peering.

5.1. IDT

IDT's customer base of 100,000 dial-up access customers in 1998 is small compared to MCI's more than 300,000 customers and far more diversified service offerings, placing relatively low

pressure on IDT's network resources. This allows IDT's small network — smallest of the ten studied in this paper — to perform very efficiently in the median download category. IDT takes advantage of its small network diameter (4) to route traffic efficiently across its network. This allows IDT to service its mostly dial-up customer base with a minimum of bandwidth and links.

The second contributor to IDT's fastest download times is the type of traffic that the network carries. IDT (International Discount Telecommunications) helped create the international "call-back" telephone service. Call-back service allows anyone in the world to bypass expensive international rates and place a call as if phoning from inside the USA. From its origins in call-back service, IDT has branched into Internet telephony, recently signing an agreement with Cable & Wireless to route C&W's telephone traffic onto IDT's network. Voice traffic utilizes significantly less bandwidth on IDT's network than data traffic. This allows IDT's network to reserve plenty of bandwidth for the small data demands of its dial-up Internet access customers, resulting in top-ranked download speeds.

5.2. CompuServe

CompuServe, on the other hand, has a mismatch between its declining customer base and its bulked-up network resources. As with other network providers, acquisitions and mergers have had a profound impact on CompuServe's network. In early 1998, WorldCom purchased CompuServe's network infrastructure from H&R Block for $1.2 billion. To complicate matters, at the same time, WorldCom purchased ANS Communications from America Online (AOL) and signed a five-year deal whereby WorldCom became AOL's primary network service provider. In exchange for the five-year contract, AOL received CompuServe's customer base, Interactive Services Division and $175 million in cash (Boardwatch, 1998, p. 233). CompuServe is slated to become WorldCom Advanced Networks and become interconnected with WorldCom's backbone. The result is a network with prodigious capacity and few customers to produce traffic to congest it.[4] This can be seen in the network rankings, where CompuServe's median download time is second overall but no higher than fourth in the graph-theoretic measures (Table 6).

In the graph-theoretic measures dealing with efficiency through connectivity and redundancy, the *alpha* index, *gamma* index and diameter, Compuserve ranks third, fourth and fourth respectively, nearly equal to its second-ranked median download time. Both IDT and CompuServe show that a small and efficient network can produce quick download times with minimum of bandwidth (IDT and Compuserve have the lowest total bandwidth totals of all the networks examined), as long as traffic is moderate.

5.3. Frontier

In the case of Frontier, the distinction between graph-theoretic network measures (cyclomatic number, *beta*, *alpha*, and *gamma* indices, and diameter) and non-graph measures (median download and total bandwidth) is especially significant. Frontier comes out on top in all of the

[4] CompuServe saw a decline during 1998 from over 5 million to 2 million customers (Telecommunications Reports International, 1999).

graph-theoretic measures, but much further down the list in the non-graph-theoretic measures. This is explained by the fact that Frontier is a fully meshed and interconnected non-planar network: every node (city) is connected by an edge (fiber optic cable) to every other node in the network. A fully connected network produces the highest and most efficient graph-theoretic numbers. This efficiency does not carry over to the non-graph-theoretical measures because the Internet cannot be understood fully using only traditional graph theory.

While Frontier has a very efficient topological network, it does not have a very efficient Internet network. Frontier serves only ten cities with (rapidly) outdated 45 Mbps DS-3 fiber optic trunks. In order to reach any of the other 50 metropolitan areas served by fiber optic backbones, it must depend on other companies' networks through public or private peering. Public peering is fraught with the peril of high packet loss, causing a high packet retransmission rate and resulting in slow transmission rates, further compounded by slow 45 Mbps connections. Frontier emphasizes "digital distribution" for Web customers, with mirrored Web sites, software, and hardware to determine the best paths for the quickest downloads. Frontier believes this will position them well for the high-bandwidth IP transit services for all applications (*Boardwatch*, 1998, p. 109). The problem is that Frontier failed to respond to Amdahl's law — that system bandwidth is set by the slowest component (Gilder, 1998). Frontier's 45 Mbps trunks suggest that the firm is poorly positioned for the high-bandwidth IP transit services, especially with other providers using or upgrading to OC-12 622 Mbps links and new-generation networks utilizing wave dimension multiplexing (WDM) technologies for OC-192 10 Gbps links.

Frontier also is limited in its potential for expansion because of the large investment required to upgrade a fully connected and redundant network. Frontier would have to upgrade 45 fiber optic links to reach only 10 metro areas, whereas other networks can upgrade fewer lines and increase service to a larger number of cities. In addition, Frontier's higher cost of upgrading a fully interconnected network would require a far greater length of fiber optic cable to connect all geographically distributed places to each other. While a fully meshed topology might be theoretically efficient, it does not functionally or economically make sense for Internet networks.

5.4. Verio

The same gap between topological efficiency and lackluster end-user download times is found with Verio, which consistently ranks in the top three or four networks in several rankings in Table 6. With 56 edges, 26 vertices and 5372 Mbps of total bandwidth, Verio is a relatively large network. The one incongruity is Verio's poor performance in *Boardwatch*'s median download testing. This poor performance points to the shortcoming of the lack of data on the size on the network's customer base. Customers create traffic, which utilizes network resources, creates congestion and has a profound impact on download speeds. While Verio's customer data are not available, a look at the firm's recent corporate and market activity tells part of the story. In 1998 alone Verio acquired several Internet service providers and Web hosting firms, resulting in a huge expansion of its customer base and network traffic.

Each of Verio's acquisitions represents a big increase in customers, increasing traffic and straining network resources. This would explain Verio's poor performance in median download measures. In March, 1998, Verio entered into a 15 y Capacity and Services Agreement with Qwest Communications Corporation, under which Verio will have access to long-haul capacity and

ancillary services on Qwest's planned 16,285 mile MacroCapacity (sm) Fiber Network. The agreement with Qwest, which will be running OC-192 10 Gbps lines, should expand Verio's network resources to match its increasing customer base and traffic. Verio continues a rapid program of expansion (Hurley, 1999).

5.5. Savvis

Companies can overcome network topological deficiencies with innovative technological solutions. This is the case with Savvis Communications Corporation. In both graph and non-graph measures, Savvis lurks towards the bottom of the rankings, yet in median download time it performs significantly better (Table 6). Savvis' technological solution for network efficiency is its dispersed private network access points (private NAPs) which, according to the company, "spatially collapses the Internet." Savvis describes the private NAP technology as a provider of direct transit connectivity to the three major Internet backbone providers (MCI, Sprint and UUNET) (Savvis, 1999). The three backbones currently represent connectivity to over 80% of worldwide Internet sites. Whenever a customer sends or receives data, the private NAP router coordinates with the Internet backbone providers' routers to find the best available path for the data (Boardwatch, 1998, p. 191). This allows customers' data to avoid the congested public NAPs and find the most efficient route across the Internet.

In conjunction with other corporate partners, Savvis is planning on marketing its private NAPs as a peering solution. It plans to deploy brokered peering points throughout the US (where peering requirements would be publicly brokered). This would take private peering out of private contracts and into the public where smaller networks can play on a level field. Whether the large networks (MCI (C&W), Sprint and UUNET) will agree to include themselves in such a proposal remains to be seen.

5.6. AT&T and IBM

Surprisingly low on the list of big players in IP networks is traditional telecom powerhouse AT&T. While other firms jumped on the Internet bandwagon, AT&T kept its focus on the long-distance telephone market. In 1997, the firm's strategy was revamped to an Internet and wireless focus, including plans to invest $2.9 billion in its IP network alone in 1999. A larger strategic move was AT&T's purchase of IBM's Global Network for $5 billion in late 1998. Although some big Internet players, such as Qwest, do not see AT&T as a threat (Elstrom, 1999), network analysis of IBM and AT&T reveals that Qwest might have underestimated its competition. While both AT&T and IBM show up only in the middle of all the network efficiency indicators (Table 6), they show up side-by-side in all network measures except median download time and diameter. A comparison of the numbers shows the similarity between the two networks (Table 7).

The topological similarity between the networks would suggest that integration of the two should be relatively easy, perhaps because parts of IBM's network were originally leased from AT&T. According to an anonymous industry source, the network integration would most likely occur at network consolidation points where network nodes will be connected by broadband fiber links. These direct connections between the two networks are where upgrades will be made first, as

Table 7
IBM and AT&T networks compared[a]

Network indicator	IBM	AT&T
Edges	22	24
Vertices	15	16
Cyclomatic number	8	9
Beta	1.467	1.5
Alpha (%)	8.79	8.57
Gamma (%)	20.95	20.00
Diameter	5	9
Median download time (seconds)	3.290	2.385
Total bandwidth (Mbps)	990	1,080

[a]Source: Calculated from data in *Boardwatch* (1998).

was done when AT&T acquired CERFnet. Further, with AT&T's purchase of @Home's OC-48 Cable Network, AT&T will have access to 15,000 miles of IP over Dense Wave Dimension Multiplexing (DWDM) fiber (the latest fiber technology). The IBM and @Home integration could prove to be another strategic advantage for AT&T's move into the IP network and data market. (This analysis does not consider IBM's international connections to 100 countries.) One area where difficulty could arise is in the interface between the IBM network's aging DS-3 frame-relay technology and the newer OC-3 ATM links in AT&T's network. The faster technology is the likely explanation for AT&T's faster download speed despite a larger network diameter.

5.7. UUNET and Sprint

UUNET and Sprint are, along with MCI (C&W), among the top three dominant Internet backbone providers. UUNET and Sprint account for 43.7% of all Internet backbone connections; the addition of MCI raises the share to 72.2% of connections, but the current state of C&W's acquisition leaves the picture unclear. This vast number of interconnections is made possible by a large number of fiber optic links, bolstered by high bandwidth per link. UUNET ranks first in number of edges and Sprint fourth but, more importantly, they are first and second in total available bandwidth (Table 6). UUNET accomplishes this by having a very large network that connects many locations with a mix of high bandwidth links and more routine links for less congested routes. This network size advantage is seen in UUNET's top-ranked cyclomatic number and third-ranked *beta* index.

Sprint, on the other hand, accomplishes its high bandwidth and connection numbers with a well-connected redundant network. Sprint connects fewer than half the nodes that UUNET connects, but provides redundant paths for 27.4% of routing possibilities, compared to 7.0% for UUNET. Further, Sprint provides more efficient routing, with 38.1% of its network being interconnected, compared to 12.9% of UUNET's network. While their paths are different, both UUNET and Sprint have gained dominant positions in the network market. With the power to deny smaller ISP access to 43.7% of the Internet, both UUNET and Sprint have been able to mandate their own peering policies.

Sprint maintains control of its peering by filtering its routing tables. "Sprint has been aggressive in deleting routes to small Class C addresses to control routing table size and as a result, not all Internet sites may be available from the Sprint network" (*Boardwatch*, 1998, p. 201). UUNET, on the other hand, controls peering by set policy measures. In 1997, UUNET mandated it would no longer accept peering requests from ISPs who cannot route traffic on a bilateral and equitable basis — that is, only with other ISPs that operate a national network with a dedicated, diversely routed DS-3 or faster backbone, which connects with UUNET at 45 Mbps or greater speed in at least four geographically diverse locations (*Boardwatch*, 1998, p. 242). Networks that cannot meet these criteria must pay monthly rates to connect to UUNET's backbone. Small and regional networks and Internet content providers are at a disadvantage in their relationship with UUNET and Sprint, but potential regulatory intervention remains unresolved. The evolving nature of hierarchical bilateral peering agreements suggests that if ISPs do not play close attention to network design, geographical diversity, and technological deployment they will pay a hefty price in today's market.

5.8. Cable & Wireless

At the other end of the interconnection spectrum is Cable & Wireless (C&W), which was the most weakly connected of the ten networks in 1998 according to the graph-theoretical measures (Table 6). C&W is at the bottom of the non-graph measures as well. While Internet-specific aspects of networks can cause networks with high graph-theoretic measures to perform poorly, there is a direct link between graph theoretical measures and network performance. Technology can limit the performance of topologically well-designed networks. C&W's network structure, not its technology, limited the network's performance. Technologically, C&W was well-provisioned with state-of-the-art equipment. The problem lay in the design of the network: a series of topological stars that impose a hierarchical structure on the network. Central nodes in the Chicago, Dallas, and Miami areas routed traffic along single links to nearby nodes. In such a star design, most nodes are connected to only one other node (the hub node), thus allowing only one routing option for each node and producing an inefficient network. The impact is clearly seen in C&W's low graph theoretical scores, especially the very low *beta*, *alpha*, and *gamma* scores, indicating low degrees of interconnectedness and redundancy (Tables 5 and 6). Redundancy and interconnectedness play vital roles in Internet efficiency, but these are not the only shortcomings of the network's design. The packet-switching technologies of the Internet depend on packets being able to take different routes to the same destination where they are reassembled. This ability provides increased efficiency and avoids network link failures. In the star topology, most node pairs have only one routing possibility for packets to take, resulting in higher packet traffic congestion and link failure.

The star topology has been abandoned by C&W in the wake of the regulatory approval of the MCI–WorldCom merger, as a result of which MCI sold its Internet backbone to Cable & Wireless. The sale of MCI's network temporarily left the new C&W network in a state of limbo. While C&W gained 22 nodes in the USA, as well as 15,000 interconnection ports, routers, switches, modems, e-mail servers, and more than 40 ongoing peering agreements from the deal (*Boardwatch* 1998, p. 77), the combination of MCI's partially meshed topology and C&W's star topology would be a technical nightmare for providing quality of service in a competitive marketplace. Note that MCI

ranked second and C&W 27th in median download times among 33 US networks in late 1998 (*Boardwatch*, 1998).

Cable & Wireless could run the two networks separately, using standard private (in this case, internal) peering to move traffic between the two. A second possibility would be to integrate the two networks functionally by combining routing tables, meshing the topologies together. While data would move more efficiently, the star topology and slower bandwidth of C&W's old network would decrease the performance of the acquired MCI network. A third possibility, that C&W would simply scrap its old network and use only the acquired MCI network, gained credibility when topologies of C&W's new network were published (in February, 1999), revealing a network that looks remarkably similar to MCI's network topology alone (McCarthy, 1999). The combined MCI–C&W in mid-1999 is a network that is now little different from the former MCI network and towards the top of all of the network measures, especially in download time and total bandwidth (Table 8).

Comparison of the pre- and post-merger topologies (Fig. 1) suggests that not only has C&W removed its old ring-linked star design, but it also has modified the former MCI network (*Boardwatch*, 1999; Cable & Wireless, 1999). After acquiring the MCI network C&W bolstered its bandwidth by installing double, triple and quadruple OC-12 links between most nodes and, at the same time, chose not to utilize many direct long-haul links between several large urban areas. The result is a network that has declined in connectivity in graph-theoretical measures (*beta*, *alpha*, and *gamma*), but increased in bandwidth (Table 8). The result is 52% reduction in median download time in comparison with C&W's old network, raising to fastest among the ten networks included here, and a slight improvement (0.7%) in comparison with the median download time of MCI's network.

This result in performance suggests that the economics of IP networks allow direct connections between two high-traffic nodes to be replaced by higher-bandwidth multiple-hop connections without performance loss. This would support the view that C&W has followed a network strategy

Table 8
Network analysis of the MCI addition to Cable & Wireless (C&W)

Network indicator	C&W (rank)[a]	C&W after acquisition of MCI's network (rank)[a]
Edges	38 (5)	33 (5)
Vertices	36 (1)	18 (4)
Cyclomatic number	3 (8)	6 (7)
Beta	1.056 (10)	1.833 (5)
Alpha (%)	0.336 (10)	11.765 (3)
Gamma (%)	6.03 (8)	21.57 (4)
Diameter	9 (8)	5 (4)
Median download time (seconds)	3.85 (10)	1.845 (2)
Total bandwidth (Mbps)	2675 (4)	39,202 (1)

[a] Ranks are relative to rankings among the ten networks shown in Table 6.
Source: Calculated from data in *Boardwatch* (1998), *Boardwatch* (1999), and Cable & Wireless (1999).

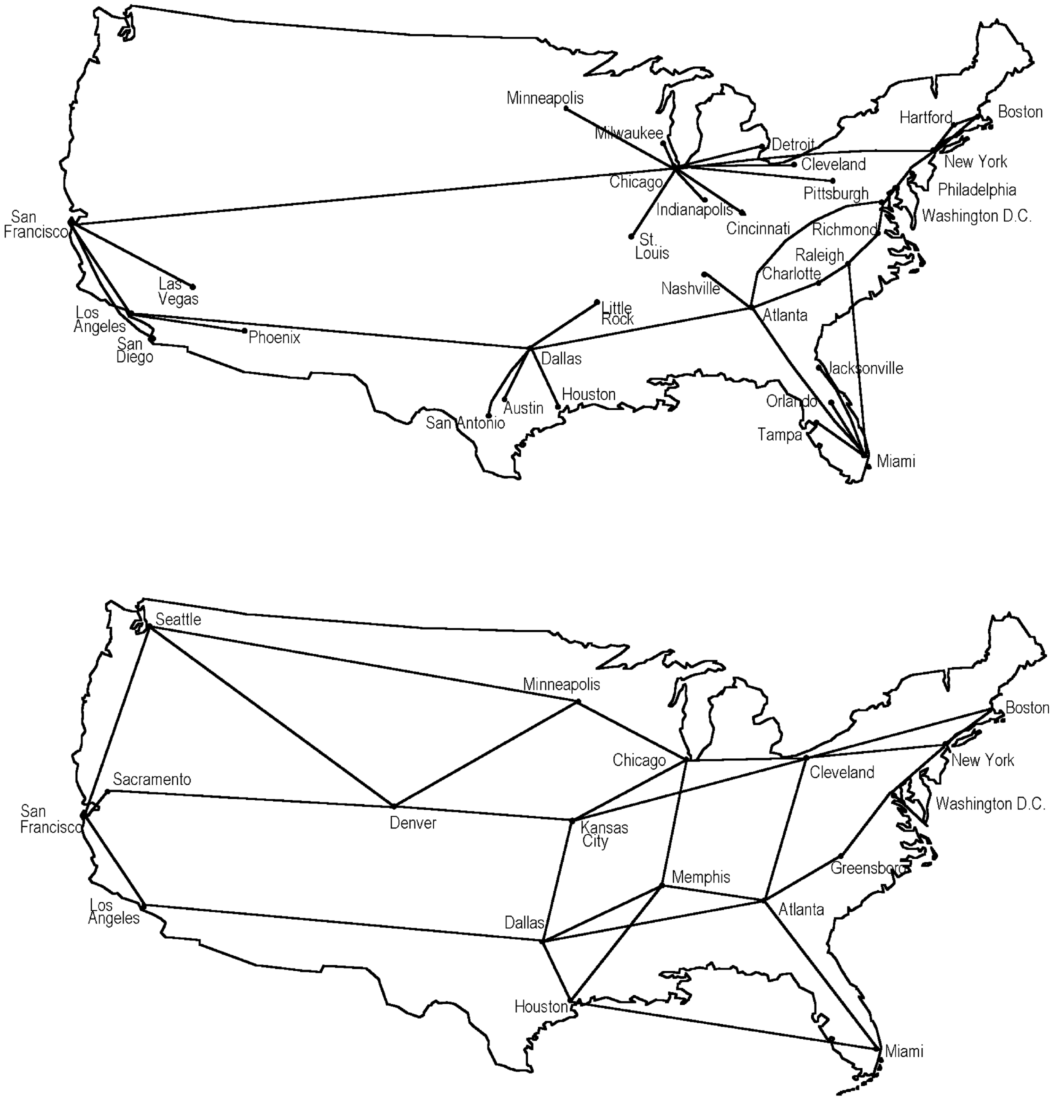

Fig. 1. Cable & Wireless USA Internet backbone network before (top) and after (bottom) acquisition of MCI backbone network.

with the premise that bandwidth is a greater determination of network performance and latency than the number of hops between connections. Other backbone providers have added bandwidth to their networks as well, but few as much so as C&W.

Equally interesting but more subtle is C&W's new geographic allocation of network resources. The US Internet has a "coastal effect," wherein cities on the eastern and western seaboards are better connected to each other than they are to the interior of the country. Functionally, this puts the majority of US coastal cities at the core of the Internet and cities located in the interior more at the periphery of the network (Malecki and Gorman, forthcoming). However, the C&W network exhibits the opposite pattern (Table 9). C&W's post-merger network was put into a connectivity

Table 9
Relative ranking of Internet connectivity of coastal and interior locations[a]

Urban area	Number of cities reachable in 1 hop	Number of cities reachable in 2 hops	DIFF	Geographical classification
San Francisco	3	8	5	Coastal
Sacramento	2	6	4	Coastal
Los Angeles	3	10	7	Coastal
Seattle	3	9	6	Coastal
Denver	3	10	7	Interior
Dallas	5	19	14	Interior
Houston	3	11	8	Interior
Atlanta	5	18	13	Interior
Memphis	4	17	13	Interior
Greensboro	2	7	5	Coastal
Miami	2	8	6	Coastal
Washington DC	2	6	4	Coastal
New York	4	12	8	Coastal
Boston	2	9	7	Coastal
Chicago	4	16	12	Interior
Minneapolis	3	10	7	Interior
Cleveland	5	19	14	Interior
Kansas City	4	17	13	Interior

[a]Source: Calculated from data in *Boardwatch* (1999) and Cable & Wireless (1999).

matrix to determine the number of direct (1-hop) connections; the matrix then was multiplied to find the number of locations that could be reached in two hops. The increase in the number of places that can be reached in two hops (over those reachable in one hop) is documented as DIFF. The DIFF column shows that in C&W's network cities located in the interior are more connected and accessible (and offer more routing options) than the cities on the coasts, the exact opposite of the pattern found in the US Internet as a whole (Malecki and Gorman, forthcoming). Further, this effect was increased by C&W's decision to remove long-haul links that circumvented connection to interior hub nodes. The decision to reduce long-haul links in the network may reflect simply a choice of greater bandwidth over hop minimization but, in any event, the change in geographic focus raises some interesting possibilities.

For example, by allocating more bandwidth to connect to the interior instead of spending resources on long-haul links that connect coast-to-coast, C&W provides interior cities with more Internet resources than the majority of its competitors. These resources are further bolstered by the choice to allocate double, triple and quad OC-12 lines into the interior. This raises questions about the effects of the interconnectivity of the Internet as a whole. Because the Internet largely follows hot potato routing (where data hops on the nearest location of the network for the long-haul transit to its destination) C&W could be exploiting this interconnective aspect of the Internet as a competitive advantage of its network design. With the 40 ongoing peering agreements that C&W acquired from MCI, C&W is able to utilize the long-haul routes of other networks to pass traffic through hot potato routing to non-C&W destinations, reducing the economic and traffic burden to

its network significantly. The use of graph theory for IP network analysis opens many interesting avenues of inquiry and investigation.

6. Conclusions

The Internet backbone market is rapidly changing the dynamic telecommunications infrastructure. Graph-theoretic network measures provide a valuable topological analysis of the efficiency of these networks, but tell only part of the story of end-user performance on the Internet. Graph theory must be utilized as a tool in conjunction with knowledge of the Internet backbone market structure, technical issues, and information about proprietary networks. The combinations of these analyses allows us to speculate about the complexities of the Internet from an economic, geographic, and technical perspective.

The network structure and analysis of the Internet, as illustrated in this paper, has implications for the urban hierarchy, business location and the future expansion of the Internet itself. In the new Internet, network providers are the dominant players, exerting their power especially in the context of peering. The backbone providers' wide-ranging domination is based on the prominence of the profit motive as the primary impetus of network growth and resource allocation. Simply put, the Internet will be located at the level of quality where money and demand puts it. That demand is concentrated in large urban areas. The Internet is not a utopian public good available to everyone, whether core or periphery and, furthermore, it is not available at the same level of technology and service to all locations (*Fortune*, 1999).

Less well-known is the impact of individual provider networks. Despite advertising by the backbones and other Internet providers, users are unable to plow through the numerous indicators of network performance (Paltridge, 1998b). Similar problems afflict Web statistics (Vonder Haar, 1999). The complexity of Internet infrastructure is apparent to users, operators, and vendors alike, although each group has different goals (Claffy & Monk, 1997).

The network analysis presented in this paper illustrates that all networks are not alike, and that they can be analyzed in other than real-time performance, which is changing constantly. The backbone networks themselves are changing as well, but the bulk of their infrastructure is a sunk cost and generally is the foundation to which additions are made.

Acknowledgements

The authors appreciate the constructive comments and suggestions of two anonymous reviewers.

References

Antonelli, C. (1997). A regulatory regime for innovation in the communications industries. *Telecommunications Policy*, *21*, 35–45.

Baccala, B. (1998). BGP topological models. *Connected: an Internet encyclopedia* < http://www.inf.tsu.ru/RFC/Connected/index.htm. >

Bailey, J. (1997). The economics of Internet interconnection agreements. In L. W. McKnight, & J. P. Bailey, *Internet economics* (pp. 155–168). Cambridge, MA: MIT Press.

Boardwatch (1998). *Boardwatch Magazine's directory of Internet service providers: guide to Internet access and the World Wide Web*. Vol III, No. 2. Golden, CO: Mecklermedia. < http://boardwatch.internet.com/isp/spring99 >

Boardwatch (1999). *Boardwatch Magazine's directory of Internet service providers* (11th ed.) Golden, CO: Penton. < http://boardwatch.internet.com/isp/summer99/ >

Bond, J. (1997). Telecommunications is dead, long live networking. *Public Policy for the Private Sector* (World Bank), *119*, 1–4.

Burch, H., & Cheswick, B. (1999). Mapping the Internet. *Computer*, *32*(4), 97–102 < computer.org/computer/IntWatch0499.htm >

Cable & Wireless USA (1999). Internet backbone. < http://www.cw-usa.net/ internet_backbone.htm >

Claffy, K., & Monk, T. (1997). What's next for Internet data analysis? Status and challenges facing the community. *Proceedings of the IEEE*, *85*(10), 1563–1571.

Claffy, K., Monk, T. E., & McRobb, D. (1999). Internet tomography. *Nature Web Matters*, 7 January < http://helix.nature.com/webmatters/tomog.html >

Cukier, K. N. (1998). The global Internet: a primer. *Telegeography 1999.* Washington, DC: Telegeography, 112–145.

Cukier, K. N. (1999). Bandwidth colonialism? The implications of Internet infrastructure on international e-commerce. *INET99 Proceedings*. Reston, VA: Internet Society < http://www.isoc.org/inet99/proceedings/1e/1e_2.htm >

Elstrom, P. (1999). Mike Armstrong's strong showing. *Business Week*, January 25, 94–96.

Evagora, A. (1997). Borne in the U.S.A. *Tele.com*, September < http://www.teledotcom.com/0997/features/tdc0997globe.side2.html >

Farrell, J., & Katz, M. L. (1998). Public policy and private investment in advanced telecommunications infrastructure. *IEEE Communications Magazine*, *36*(7), 87–92.

Fortune (1999). The need for speed. *Fortune Technology Buyer's Guide*, 122–126.

Gabel, D. (1996). Private telecommunications networks: an historical perspective. In E. Noam and A. Nishuilleabhain, *Private networks public objectives*. pp. 35–49, Amsterdam: Elsevier.

Gilder, G. (1998). Bandwidth debate: Worldcom and Qwest. *Gilder Technology Report 3*(3) < http://www.gildertech.com/pdf/Mar98.pdf >

Gillespie, A., & Robins, K. (1989). Geographical inequalities: the spatial bias of the new communications technologies. *Journal of Communications*, *39*(3), 7–18.

Gillett, S. E., & Kapor, M. (1997). The self-governing Internet: coordination by design. In B. Kahin, & J. H. Keller, *Coordinating the Internet* (pp. 3–38). Cambridge, MA: MIT Press.

Graham, S., & Marvin, S. (1996). *Telecommunications and the city: electronic spaces, urban places*. London: Routledge.

Haggett, P., & Chorley, R. (1969). *Network analysis in geography*. New York: St. Martins Press.

Harary, F., Norman, R. Z., & Cartwright, D. (1965). *Structural models: an introduction to the theory of directed graphs*. New York: John Wiley.

Hepworth, M. E. (1990). *Geography of the information economy*. London: Belhaven.

Huitema, C. (1995). *Routing in the Internet*. Englewood Cliffs, NJ: Prentice Hall.

Huitema, C. (1997). *IPv6: the new Internet protocol* (2nd ed.) Englewood Cliffs, NJ: Prentice Hall.

Hurley, H. (1999). Veritable Verio. *Telephony*, *236*(7), 30–38.

Huston, G. (1999). Interconnection, peering, and settlements. *INET99 Proceedings*. Reston, VA: Internet Society < http://www.isoc.org/inet99/proceedings/1e/1e_1.htm >

ICANN [Internet Corporation for Assigned Names and Numbers] (1999). *About ICANN.* < http://www.icann.com/abouticann.html >

Kansky, K. (1963). *Structure of transportation networks: relationships between network geometry and regional characteristics*. University of Chicago, Department of Geography, Research Paper 84.

Kellerman, A. (1993). *The geography of telecommunications*. London: Belhaven.

Kellerman, A. (1997). Fusions of information types, media and operators, and continued American leadership in telecommunications. *Telecommunications Policy*, *21*, 553–564.

Leo, E., & Huber, P. (1997). The incidental, accidental deregulation of data … and everything else. *Industrial and Corporate Change*, *6*, 807–828.

Lipsey, R. G., Bekar, C., & Carlaw, K. (1998). What requires explanation? In E. Helpman *General purpose technologies and economic growth*. pp. 15–54. Cambridge, MA: MIT Press.

Mackie-Mason, J. K., & Varian, H. R. (1997). Economic FAQs about the Internet. In L. W. McKnight, & J. P. Bailey, *Internet economics* (pp. 27–62). Cambridge, MA: MIT Press.

Malecki, E. J. & Gorman, S. P. (forthcoming) Maybe the death of distance, but not the end of geography: The Internet as a network. In S. D. Brunn & T. R. Leinbach, *The wired worlds of electronic commerce*.

McCarthy, B. (1999). Darth Vader behind beheadings at Cable & Wireless? *Boardwatch Magazine*, February < http://boardwatch.internet.com/mag/99/feb/bwm1.html >

Moss, M. (1998). Technologies and cities. *Cityscape*, *3*, 107–127.

Nyusten, J. D., & Dacey, M. F. (1961). A graph theory interpretation of nodal regions. *Papers and Proceedings of the Regional Science Association*, *7*, 29–42.

Paltridge, S. (1998a) *Internet traffic exchange: developments and policy*. Paris: Organisation for Economic Co-operation and Development. < http:www.oecd.org/dsti/sti/it/cm/prod/traffic.htm >

Paltridge, S. (1998b). *Internet infrastructure indicators*. Paris: Organisation for Economic Co-operation and Development < http://www.oecd.org/dsti/sti/it/cm/prod/tisp98-7e.htm >

Paltridge, S. (1999). *Building infrastructure capacity for electronic commerce: leased lines developments and pricing*. Paris: Organsiation for Economic Co-operation and Development < http://www.oecd.org/dsti/sti/it/cm/prod/traffic.htm >

Paxson, V. (1997). End-to-end routing behavior in the Internet. *IEEE Transactions on Networking*, *5*, 601–615.

Rickard, J. (1998). High noon with the Telcos, FCC and ISP's. *Boardwatch Magazine*, October < http://boardwatch.internet.com/mag/98/oct/bwm1.html >

Rogers, J. D. (1998). Internetworking and the politics of science: NSFNET in Internet history. *The Information Society*, *14*, 213–228.

Salomon, I. (1996). Telecommunications, cities and technological opportunism. *Annals of Regional Science*, *30*, 75–89.

Savvis (1999). Savvis Communications Corporation Web site < http://www.savvis.com >

Schiller, D. (1999). *Digital capitalism: networking the global market system*. Cambridge, MA: MIT Press.

Srinagesh, P. (1997). Internet cost structures and interconnection agreements. In L. W. McKnight, & J. P. Bailey, *Internet economics* (pp. 121–154). Cambridge, MA: MIT Press.

Taaffe, E. J., Gauthier, H. L, & O'Kelly, M. E. (1996). *Geography of transportation* (2nd ed.) Upper Saddle River, NJ: Prentice Hall.

Telecommunications Reports International (1999). Telecommunications Reports International's Top ISPs by Consumer Subscriber Base. *Boardwatch Magazine's directory of Internet service providers* (11th ed.) Golden, CO: Penton, 9. < http://boardwatch.internet.com/isp/summer99/topisps.html >

Vonder Haar, S. (1999). Web metrics: go figure. *Business 2.0*, *4*(6), 46–47.

Wan, H. A., & Chung, C. (1998). Web page design and network analysis. *Internet Research*, *8*, 115–122.

Wheeler, D. C., & O'Kelly, M. E. (1999). Network topology and city accessibility of the commercial Internet. *Professional Geographer*, *51*, 327–339.